IN THE UNITED STATES DISTRICT COURT
SOUTHERN DISTRICT OF NEW YORK

RELX INC., Plaintiff, vs. INFORMATICA CORP., Defendant.	1:16-cv-9718-AKH
INFORMATICA LLC, Counterclaim Plaintiff, vs. RELX INC. and RELX GROUP PLC, Counterclaim Defendants.	

EXPERT REPORT OF CHRISTOPHER RUCINSKI

Table of Contents

CONFIDENTIAL

I. INTRODUCTION

1. My name is Christopher Rucinski. I am a Vice President at Stroz Friedberg, an Aon Company ("Stroz Friedberg"), located at 53 State Street, Boston MA 02109. I have been retained by counsel for RELX, Inc. and RELX Group PLC (collectively, "RELX") to offer my opinions as an expert in computer science related to the above captioned litigation between RELX and both Informatica Corp. and Informatica LLC (collectively, "Informatica").

2. This expert report ("Report") summarizes my current opinions, which are subject to change depending on ongoing discovery and additional information. This Report contains confidential information, including information contained in documents identified by Bates numbers that were produced by the parties to this litigation during discovery. This Report is subject to the protective order agreed to by Informatica and RELX on May 26, 2017.

II. QUALIFICATIONS AND COMPENSATION

3. I graduated *cum laude* with an A.B. in Computer Science from Princeton University in 2010, and in 2015, I obtained the GCFE (GIAC[1] Certified Forensic Examiner) certification. In 2010, I began working at Elysium Digital as a computer scientist, and I continued to work there until Elysium Digital was acquired in 2015 by Stroz Friedberg, where I have worked since. At Elysium Digital and Stroz Friedberg I have consulted on more than 80 technical matters to date, most of which have involved some form of technical document review.

4. I have provided consulting services for the Federal Trade Commission, and I provided deposition testimony in 2013 in the matter of *Shurtape Technologies, LLC et al. v. 3M Company*, U.S. District Court of Western North Carolina (Case No. 5:11-cv-00017). I also provided deposition and trial testimony in 2015 in the matter of *BMG Rights Management (US)*

[1] "GIAC" stands for Global Information Assurance Certification: http://www.giac.org.

LLC et al. v. Cox Enterprises, Inc. et al., U.S. District Court, Eastern District Court of Virginia (Case No. 1:14-cv-0161).

5. My CV is attached as Exhibit A.

6. Stroz Friedberg is being compensated for my time at a rate of $785 per hour. Stroz Friedberg is also being compensated at rates of $265 to $785 per hour for the work of additional Stroz Friedberg employees who are working at my direction. All compensation described above does not depend on the outcome of this case.

III. MATERIALS CONSIDERED

7. I considered the following materials in arriving at my opinions in this Report:

i. Dkt. 1 ("Complaint"), Dkt. 1-1 ("Reihl Declaration"), Dkt. 17, Dkt. 18 ("Counterclaim"), Dkt. 19, Dkt. 20, Dkt. 21 ("Billingsley Declaration"), Dkt. 22, Dkt. 23, Dkt. 24, Dkt. 65, Dkt. 66, Dkt. 67, Dkt. 68, and Dkt. 72 ("Mishra Declaration")

ii. INFA_0000218566-INFA_0000218570, INFA_0000218583, INFA_0000230313, INFA_0000243617, INFA_0000255246, INFA_0000255255, INFA_0000259688, INFA_0000259793, MISHRA0000022029-MISHRA0000022032, RELX044494, RELX050982-RELX050986, RELX050991-RELX050994, RELX054590-RELX054592, RELX054593-RELX054595, RELX054603-RELX054606, RELX054607-RELX054610, RELX054611-RELX054614, RELX054618-RELX054622, RELX054623-RELX054627, RELX054628-RELX054632, RELX086341, RELX086342, RELX086343, RELX152655, RELX204919-RELX204920, RELX206412-RELX206524, RELX207230-RELX207235, RELX207236-

RELX207242, RELX234894, RELX263479-RELX263480, RELX263481-RELX263483, RELX263484, RELX263485-RELX263487, RELX263488-RELX263491, RELX263492-RELX263495, RELX263496-RELX263499, RELX265066, RELX270761-RELX270770, RELX271259, RELX271261, RELX271262-RELX271263, RELX271264-RELX271265, RELX271273-RELX271274, RELX271287-RELX271288, RELX281935, RELX283550, RELX283551, RELX283552, RELX283553, RELX283931-RELX283932, RELX283933-RELX283950, RELX283951-RELX283962, RELX283963-RELX284023, RELX284024, RELX284025-RELX284034, RELX284035-RELX284036, RELX284037-RELX284054, RELX284055-RELX284115, RELX284116, RELX284117-RELX284126, RELX284129-RELX284130, RELX284131, RELX284132-RELX284192, RELX284193-RELX284202, RELX284203-RELX284220, RELX311858-RELX311863, RELX311871-RELX311877, RELX311878-RELX311884, RELX314107-RELX314109, RELX318833, RELX346794-RELX346796, RELX346802-RELX346805, RELX346823-RELX346829, RELX346844-RELX346851, RELX349058-RELX349059, RELX349060-RELX349062, RELX349069-RELX349071, RELX349072-RELX349075, RELX353805-RELX353816, RELX353875-RELX353886, RELX366034-RELX366036, RELX366622-RELX366623, RELX409196-RELX409198, RELX409199-RELX409200, RELX409201-RELX409203, RELX409540-RELX409541, RELX409542-RELX409543, RELX409544-RELX409545, RELX409546-RELX409548, RELX409549-RELX409551, RELX409552-RELX409553, RELX409554-RELX409556,

RELX409557-RELX409559, RELX409560-RELX409562, RELX409563-RELX409565, RELX409566-RELX409569, RELX409600-RELX409602, RELX411147-RELX411148, RELX411149-RELX411151, RELX411152-RELX411154, RELX411155-RELX411157, RELX411158-RELX411161, RELX436217, RELX436798, and RELX436799-RELX436843

iii. RELX223409-RELX223728 ("Informatica Administrator Guide Version 9.6.1")

iv. Expert Discovery000002 ("ICCE Files Processed Report")

v. Expert Discovery000008 ("ICCE Server Utilization Report")

vi. Expert Discovery000010 ("ICCE Processed Documents Percentage Report")

vii. Expert Discovery000011 ("PC Upgrade_First_Node_Prod.sh"), Expert Discovery000012 ("Upgrade Steps.txt"), Expert Discovery000013 ("Upgrade_Each_Node_Prod.sh"), and Expert Discovery000020 ("Install.sh"), collectively "Upgrade Scripts"

viii. Expert Discovery000001, Expert Discovery000003-Expert Discovery000007, Expert Discovery000009, Expert Discovery000014-Expert Discovery000019

ix. Declaration of Dwight Groff dated May 12, 2017

x. Declaration of Catherine W. Appel dated May 13, 2017

xi. Deposition of Nalin Mishra dated May 31, 2017

xii. Declaration of Dwight Groff dated August 4, 2017

xiii. Deposition of Dwight Groff dated January 17, 2018

xiv. Phone conversations with Dwight Groff and Jeffrey Hoffman of RELX

xv. Any documents, devices, software, or articles cited herein

IV. OVERVIEW OF THE LICENSING DISPUTE BETWEEN INFORMATICA AND RELX

8. From May 2010 through November 2017, RELX and Informatica entered into several licensing agreements wherein RELX, of which LexisNexis is a subsidiary, was entitled to use at least Informatica's B2B Data Exchange product, Informatica's B2B Data Transformation product, and Informatica's PowerCenter Enterprise Grid product (collectively, "Informatica Software"). The licensing agreements allow for Informatica Software products to be used on a certain number of "CPU cores," where each license purchased by RELX entitled RELX to run Informatica Software components on one CPU core.[2] The first of these license agreements, entitled "Master Software License Agreement" (hereinafter "MSLA") was executed on May 28, 2010, and includes licenses for 16 "B2B Data Exchange Production CPU cores."[3] My understanding is that RELX and Informatica dispute the number of licenses to which RELX was entitled over various times during the Relevant Time Period (defined below).[4] Informatica alleges that the Informatica Software was deployed on RELX computers with collectively more cores than RELX licensed during certain portions of the Relevant Time Period and that RELX benefited from the alleged over-deployment of the Informatica Software.[5]

V. SUMMARY OF OPINIONS

9. My opinions are based on my review and analysis of the materials listed above and cited in this Report. My opinions are described in detail in the subsequent sections of this Report, and my opinions can be summarized as follows:

[2] "CPU" stands for "Central Processing Unit" and is the component of a computer responsible for executing instructions from computer programs.

[3] See, e.g., Complaint, Exhibit A.

[4] See, e.g., Complaint, paragraphs 19-20 and Counterclaim, paragraphs 19-20.

[5] See, e.g., Counterclaim, paragraphs 44 and 71.

10. The Informatica Software is only one component of a complex system that provides the functionality of RELX's Lexis Advance® and Lexis.com® products.

11. During the Relevant Time Period, the Informatica Software was involved in processing a limited percentage of files that were made available to RELX customers via the Lexis Advance® and Lexis.com® products.

12. The technical design and reporting of Informatica's licensing arrangement with RELX failed to include several easily implemented components that would have made compliance easier for RELX to maintain.

13. Informatica's statements of work with RELX are most reasonably interpreted to require Informatica to not over-deploy the Informatica Software.

14. Informatica's license management report is confusing and uses incorrect metrics to report license compliance.

15. The documentation for Informatica's license management report is confusing and contradictory.

16. Informatica failed to provide RELX with a way to check its license compliance that was consistent with the manner in which Informatica enforced RELX's license compliance.

17. Informatica's license implementation failed to automatically prevent over-deployment even though it could have easily done so and doing so would have greatly helped to prevent over-deployment.

18. RELX accrued no benefit from the installation of the Informatica Software on computers having a collective number of CPU cores greater than RELX had licensed.

19. Informatica's Sizing Model for RELX calculates at most 12 CPU cores as the appropriate number for RELX's needs, far fewer than the amount Informatica sold to RELX and

more than double what RELX is known to have utilized on average during the Relevant Time Period.

20. The ICCE Server Utilization Report and the ICCE Files Processed Report confirm that RELX never needed more than 56 CPU core licenses to meet its needs during the Relevant Time Period.

21. Historical utilization data confirms that RELX never needed more than 56 CPU core licenses to meet its needs during the late portion of the Relevant Time Period during which RELX had 56 cores deployed.

22. Historical utilization data confirms that RELX never needed more than 56 CPU core licenses to meet its needs during the middle portion of the Relevant Time Period during which RELX had 104 cores deployed.

23. Historical utilization data and historical file processing data indicate that RELX never needed more than 72 CPU core licenses during the early portion of the Relevant Time Period during which RELX had 104 or fewer cores deployed.

24. Informatica's binary executable code that RELX copied in its normal operation is different from the computer source code for which Informatica registered copyrights.

25. Informatica's certified copyright deposits contain the contents of computer source code files and do not contain the contents of any binary executable files.

26. RELX did not copy Informatica's copyrighted computer source code during the execution of the Informatica Software on RELX systems.

27. RELX did not copy Informatica's copyrighted computer source code during the upgrade of the Informatica Software in 2015.

VI. EXPLANATION OF HOW THE INFORMATICA SOFTWARE WAS USED IN RELX'S LEXIS ADVANCE® AND LEXIS.COM® PRODUCTS

A. The Informatica Software is Only One Component of a Complex System that Provides the Functionality of RELX's Lexis Advance® and Lexis.com® Products

28. RELX's Lexis Advance® and Lexis.com® products aggregate data from various sources and sell subscriptions to customers to access that data. Lexis Advance®, the more recently developed of the two products, is specifically tailored to legal professionals and focuses on providing access to information including but not limited to statutes, dockets, and verdicts.[6] Lexis.com® also provides access to documents that legal professionals would likely find useful and in addition provides access to documents from various industries as well as public records.[7] Most of the data available on Lexis.com® is also available on Lexis Advance®.[8]

29. The following steps summarize the overall process by which information from various sources makes its way to RELX's customers of the Lexis Advance® and Lexis.com® products.

i. [redacted]

[6] See, e.g., https://www.lexisnexis.com/pdf/lexis-advance/la-overview-brochure.pdf, attached to this Report as Exhibit B.

[7] See, e.g., http://www.lexisnexis.com/literature/pdfs/LO11542-0.pdf, attached to this Report as Exhibit C.

[8] This is according to conversations with Dwight Groff.

iv. [redacted]. Before 2010, RELX had implemented its own custom software, including a collection of Perl[9] and Java[10] programs, to perform the function of Step Two. Leading up to 2010, RELX decided to begin to migrate the various components of Step Two into a more organized framework and launched its Integrated Collection Conversion Environment ("ICCE Platform," pronounced "ice") project to begin that migration. After considering several options for third-party software to potentially use in the ICCE Platform, RELX decided to use the Informatica Software as part of the ICCE Platform to manage its generic process flow. In the ICCE Platform version 1.0, RELX used the Informatica Software to implement some of the specific components in the process flow, and RELX also implemented some components of the process flow themselves. The ICCE Platform was slowly expanded over time to process a subset of documents, including documents related to case law, statutes, treaties, and other documents related to law,[11] through Step Two; there were other document types that it did not process.[12]

[9] Perl is a computer programming language that is free to use and is specifically designed for text processing.

[10] Java is a computer programming language that is free to use and is better suited for larger-scale computer programs.

[11] See, e.g., Expert Discovery000003.

[12] This is according to conversations with Dwight Groff.

31. The ICCE Platform version 1.0 began processing documents on November 8, 2012, according to records provided by RELX in the ICCE Files Processed Report.[13] Over time, and in particular when it started using the ICCE Platform version 2.0 in the middle of 2014, RELX removed many specific components in the process flow that had previously been implemented using Informatica Software and instead developed its own custom components. In early November 2017, RELX stopped using the ICCE Platform to process any of the files in Step Two. I will refer to the period over which the ICCE Platform was used to process documents as the "Relevant Time Period." RELX developed a new platform called Rocket that replaced the ICCE Platform completely by November 2017. RELX built Rocket in part on a third-party product called [redacted].[14]

32. The ICCE Platform generally functions as follows in both version 1.0 and 2.0. Informatica's B2B Data Exchange component is activated when it receives a new file from a data source; it retrieves the file and certain metadata associated with the file that allows it to route the file information to the appropriate account in Informatica's Power Center Enterprise Grid product. The appropriate account is associated with an initial Power Center Workflow ("Workflow") that is activated to begin processing the file. The file is then processed through [redacted] and validate the conversion. [redacted] The Informatica Software is part of the framework for the ICCE

[13] Note that this document's filename is ICCEFileCounts.xlsx.

[14] See, e.g., https://www.talend.com/. This is according to conversations with Dwight Groff.

Platform, and the Stages of work operate within that framework to process files. Some of the Stages comprise Informatica B2B Data Transformation Stages while others comprise custom RELX implementations.[15]

33. The ICCE Platform version 1.0 and version 2.0 functioned similarly, but the Stages they implemented were slightly different, and by the time the ICCE Platform version 2.0 was implemented, more of the Informatica B2B Data Transformation Stages had been replaced with custom-built RELX Stages. The figures below from Expert Discovery000018[16] illustrate the various stages involved in the ICCE Platform versions 1.0 and 2.0. Stages marked as red boxes with "DT" were implemented with Informatica's Data Transformation software, Stages marked as blue boxes with a "J" or green boxes with a "P" were implemented with custom RELX software, and Stages marked as white boxes with no text were not used.

[15] This is according to conversations with Dwight Groff and Expert Discovery000018.

[16] This is a PowerPoint document with a filename of "ICCE_Workflow_Overviews.ppt" produced by RELX.

ICCE 1.0 Processing Workflow Overview

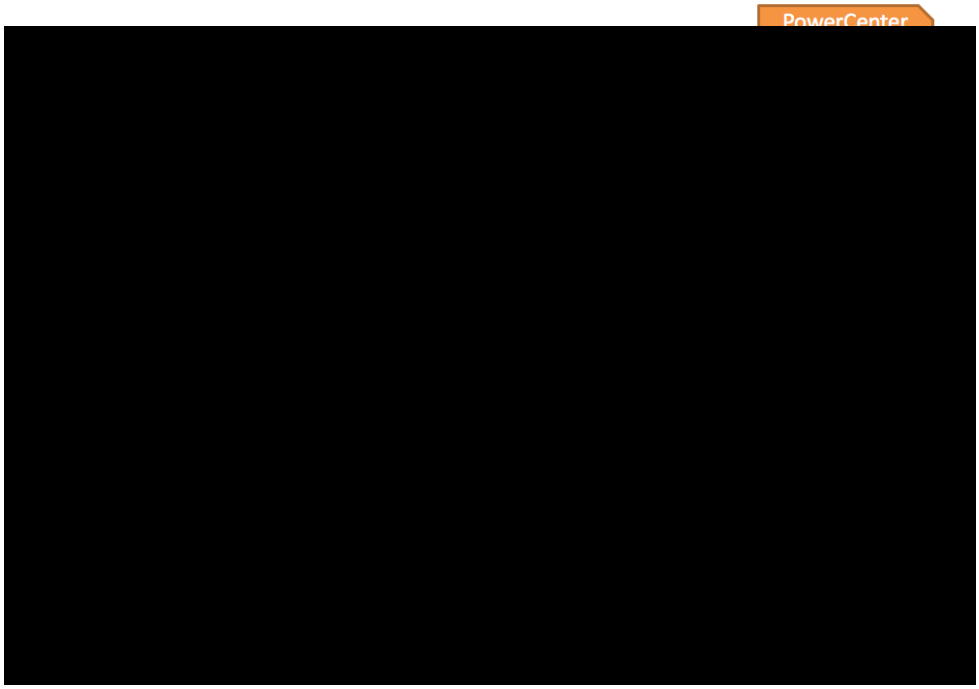

Fig. 1

ICCE 1.0 Processing Stage Technology

Fig. 2

Fig. 3

* - Workflow using DT

Fig. 4

B. During the Relevant Time Period, the Informatica Software was Involved in Processing a Limited Percentage of Files that were Made Available to RELX Customers via the Lexis Advance® and Lexis.com® Products

34. The ICCE Processed Documents Percentage Report produced by RELX is a spreadsheet that shows for each day from November 20, 2013 through November 8, 2017 the total number of documents processed through the ICCE Platform as well as the total number of documents available on the Lexis Advance® and Lexis.com® Products. Below is a graph that shows for each day the percentage of documents available on the Lexis Advance® and Lexis.com® Products that were processed through the ICCE Platform as well as a linear trend line and the average percentage across the time period covered by the ICCE Processed Documents Percentage Report, which is approximately 6.29%. The maximum percentage of documents available on the Lexis Advance® and Lexis.com® products that were processed by the ICCE Platform according to the ICCE Processed Documents Percentage Report is 15.05% on September 6, 2017.

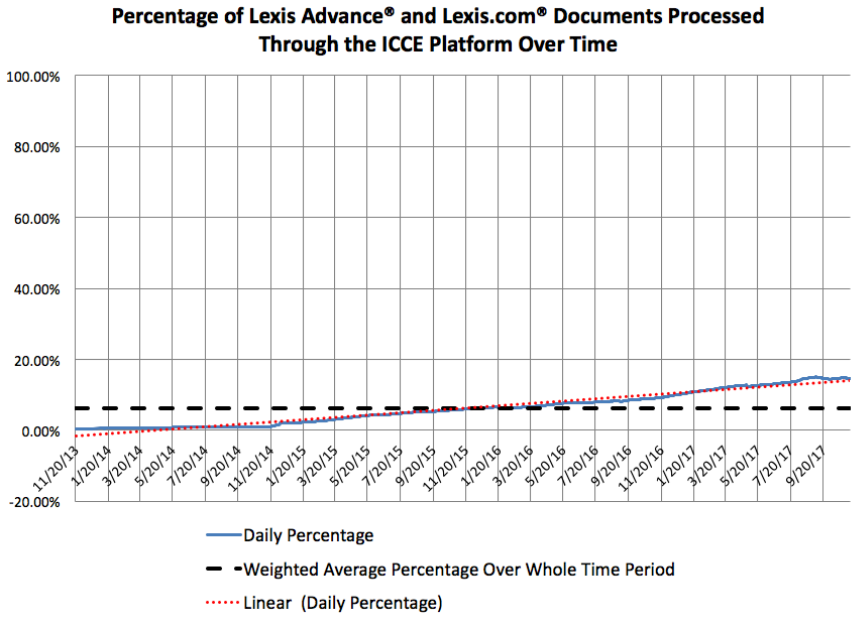

Fig. 5

35. The period of time covered by the ICCE Processed Documents Percentage Report does not cover the portion of the Relevant Time Period prior to November 20, 2013, but the positive linear trend line suggests that the daily percentages for that prior time period are generally less than the percentages for the days covered in the ICCE Processed Documents Percentage Report.

VII. THE TECHNICAL DESIGN AND REPORTING OF INFORMATICA'S LICENSING ARRANGEMENT WITH RELX FAILED TO INCLUDE SEVERAL EASILY IMPLEMENTED COMPONENTS THAT WOULD HAVE MADE COMPLIANCE EASIER FOR RELX TO MAINTAIN

A. Informatica's Statements of Work with RELX are Most Reasonably Interpreted to Require Informatica to not Over-Deploy the Informatica Software

36. Informatica and RELX entered into several Statements of Work ("SOWs") related to the Informatica Software wherein RELX agreed to pay Informatica to perform certain services. For example, Exhibit H and Exhibit J to the Complaint are two such SOWs ("Proper Use SOWs") with effective dates of December 17, 2012 and February 1, 2014, respectively. The Proper Use SOWs both identify Nalin Mishra as an Informatica employee who will perform the services listed in the Proper Use SOWs. The Proper Use SOWs both list the same services to be performed by Informatica under the heading "Design and Architecture" (p. 1 of both Proper Use SOWs):

> Design and Architecture
> - Guidance and recommendations for the design and architecture as it relates to the Informatica Platform
> - Recommendation of Informatica best practices in development and design
> - Validation of proper use of Informatica in solution design

37. The first two bullets under the "Design and Architecture" section of the Proper Use SOWs address the design and development of the Informatica Software into the RELX environment from a software engineering perspective. The first bullet under "Design and Architecture" addresses the overall design and architecture; it means that Informatica will help to

design an implementation of their software in the RELX environment. The second bullet under "Design and Architecture" addresses best practices in development and design; it means that Informatica will advise how to best develop the software that will use the Informatica Software as part of the architecture design addressed in the first bullet.

38. Though the various responsibilities for license-related compliance are not clearly stated, the Proper Use SOWs are most reasonably interpreted to require Informatica to not over-deploy the Informatica Software. The third bullet under "Design and Architecture" addresses the "validation of proper use of Informatica in solution design." The third bullet does not address RELX users' use of the software developed in the third bullet because the scope of the third bullet is limited to "in solution *design*," and the design phase of software development occurs before the software can be accessed by users.[17] The third bullet is also not reasonably interpreted to address the best ways to design the use of the Informatica Software from a software engineering perspective because the design and implementation of the Informatica Software is covered in bullet one and bullet two under "Design and Architecture." The third bullet does not explicitly mention "validation" and "proper use" in a licensing context. However, the phrases "validation" and "proper use" are not usually used together to describe software engineering implementations because there are usually several viable ways to design and implement software, each with respective benefits and costs. There are instead "best practices" with respect to software engineering, which is addressed in the second bullet, but strict dichotomies of "valid" and "invalid" or "proper" and "improper" do not usually pertain to the manner in which software is designed or implemented from a software engineering perspective. Therefore, the third bullet is

[17] Here I mean "users" in the sense of regular employees, not software engineers, accessing and using software as part of their day-to-day work.

most reasonably interpreted to mean that Informatica would not over-deploy the Informatica Software in excess of the licenses to which RELX was entitled.

39. I have seen no evidence of any way in which Informatica sought to proactively validate the proper use of the Informatica Software by RELX, as defined in the Proper Use SOWs, before the licensing audit in late 2016, which was a catalyst for this litigation, even though Informatica could have done so in a number of different ways. Nalin Mishra is listed in the Proper Use SOWs as an Informatica employee who will validate proper use, but I understand that he has disclaimed responsibility for checking whether the Informatica Software was implemented in a way that was compliant with the number of licenses that RELX purchased from Informatica.[18] Even if Nalin Mishra's responsibilities did not officially include validating that the Informatica Software he was installing for RELX complied with the number of licenses to which RELX was entitled, Informatica could have had other Informatica employees validate the proper use of the Informatica Software, in particular immediately after occasions when Informatica employees assisted in adding servers with more cores to RELX's deployment, the exact time at which a potential over-deployment would have occurred. Informatica could have also more frequently initiated the same licensing audit process that they ultimately initiated in late 2016, roughly six years after the MSLA was signed. Finally, as discussed below, Informatica could have implemented a more automated way to monitor and check for license compliance that would have reduced the need for an on-site check.

[18] See, e.g., Deposition of Nalin Mishra, May 31, 2017, 51:17 – 24.

B. Informatica's License Management Report is Confusing and Uses Incorrect Metrics to Report License Compliance

40. The Informatica Software provides the capability for a RELX administrator to automatically generate a License Management Report ("LMR") to monitor compliance with the MSLA and subsequent amendments (collectively, "License Agreement") between Informatica and RELX. The Informatica Administrator Guide Version 9.6.1 ("IAG") states this plainly at page 180 with the first sentence under the "License Management Report" heading: "You can monitor the list of software options purchased with a license and the number of times a license exceeds usage limits." Such a License Management Report ("2016 LMR") was produced as Exhibit 25 of Dwight Groff's deposition on January 17, 2018 and was generated on May 15, 2016.[19]

41. The 2016 LMR uses incorrect metrics to report license compliance because it reports license compliance according to logical CPUs instead of CPU cores as required by the License Agreement. Under the "CPU" column heading of the "Licensing" section, the 2016 LMR reports 504. According to the IAG, this number is the "Maximum number of authorized *logical* CPUs" (emphasis added).[20] However, at no point did RELX purchase 504 logical CPU licenses or even 504 CPU core licenses for the Informatica Software. Under the "Current Usage" and "Peak Usage" column headings of the "CPU Summary" section of the "Usage" section, the 2016 LMR reports 112 and 208 respectively. According to the IAG, these numbers reflect the "Maximum number of *logical* CPUs used concurrently on the day the report runs" and the "Maximum number of *logical* CPUs used concurrently during the last 12 months" respectively (emphasis

[19] See, e.g., Deposition of Dwight Groff, January 17, 2018, 389:23 – 390:4.

[20] See, e.g., IAG, page 181.

added).[21] Finally, under the "Days Exceeded License Limit" column heading of the "CPU Summary" section of the "Usage" section the 2016 LMR reports zero. According to the IAG, this number is the "Number of days that the CPU usage exceeded the license limits. The domain exceeds the CPU license limit when the number of concurrent logical CPUs exceeds the number of authorized CPUs."[22] The IAG also states at page 180:

> CPU usage. Shows the number of logical CPUs used to run application services in the domain. The License Management Report counts logical CPUs instead of physical CPUs for license enforcement. If the number of logical CPUs exceeds the number of authorized CPUs, then the License Management Report shows that the domain exceeded the CPU limit.

The "Days Exceeded License Limit" calculated in the LMR is therefore incorrectly based on logical CPUs. If RELX had indeed exceeded the license limit for zero days as reported by the 2016 LMR, then this dispute between Informatica and RELX would have no basis at all.

42. The 2016 LMR also contains a section entitled "Hardware Configuration," which lists certain properties of the computer systems used in the RELX environment, but it does so in a confusing way. It lists for each computer in the RELX environment "Logical CPU" and "Cores," which the IAG describes at page 184 as "Number of logical CPUs used to run application services in the domain" and "Number of cores used to run application services in the domain" respectively. The 2016 LMR does not list "CPU Cores" as a heading even though that is the specific language used in the License Agreement. The 2016 LMR also lists "Cores" even though the number that appears under that heading is not used in the determination of compliance with licenses, which is the ultimate purpose of the LMR.

[21] See, e.g., IAG, page 182.

[22] See, e.g., IAG, page 182

C. The Documentation for Informatica's License Management Report is Confusing and Contradictory

43. As discussed in the previous section, the LMR calculates the "Days Exceeded License Limit" on the basis of logical CPUs and not CPU cores as required by the License Agreement. The IAG, which provides documentation for the LMR, also contains confusing and contradictory statements about how licensing compliance should be determined. At pages 181 - 182, the IAG states:

> The License Management Report determines the number of logical CPUs based on the number of processors, cores, and threads.
> …
> Note: Although the License Management Report includes threads in the calculation of logical CPUs, Informatica license compliance is based on the number of physical cores, not threads. To be compliant, the number of physical cores must be less than or equal to the maximum number of licensed CPUs. If the License Management Report shows that you have exceeded the license limit but the number of physical cores is less than or equal to the maximum number of licensed CPUs, you can ignore the message.

Essentially the IAG states that the LMR incorrectly determines license compliance, which is the purpose of the LMR. Informatica could have made the LMR less confusing by clearly indicating on the LMR that the LMR did not correctly determine license compliance rather than stating that 182 pages into the IAG.

D. Informatica Failed to Provide RELX with a Way to Check its License Compliance that was Consistent with the Manner in Which Informatica Enforced RELX's License Compliance

44. In Informatica's correspondence ("2016 Correspondence") with RELX on October 7, 2016, attached as Exhibit M to the Complaint, Informatica cites a "Compliance Findings Table," which differs materially from the License Management Report discussed above, as the primary justification for Informatica's claim that RELX was not in compliance with the License Agreement. In the 2016 Correspondence, Informatica does not cite to any LMR. The Compliance Findings Table differs materially from the 2016 LMR in that unlike the 2016 LMR,

the Compliance Findings Table does not measure license compliance in terms of logical CPU cores and instead measures license compliance in terms of CPU cores, which is how licenses are also measured in the License Agreement. The Compliance Findings Table also differs materially from the 2016 LMR in that unlike the 2016 LMR, the Compliance Findings Table lists the number of physical CPU core licenses allocated to RELX, a critical fact in determining license compliance. To my knowledge, RELX only had access to the LMR and could not on its own generate a Compliance Findings Table against which RELX could verify its license compliance. Informatica therefore failed to provide RELX with a way to check its license compliance that was consistent with the manner in which Informatica enforced RELX's license compliance.

E. Informatica's License Implementation Failed to Automatically Prevent Over-Deployment Even Though It Could Have Easily Done So and Doing So Would Have Greatly Helped to Prevent Over-Deployment

45. The Informatica Software was deployed with a single license key allocated to RELX,[23] and neither that license key nor the Informatica Software in general automatically enforced the limitations of the License Agreement with respect to CPU cores even though it could have easily done so, and doing so would have greatly helped to prevent over-deployment.

46. Informatica could have easily automatically enforced the CPU core limitations of the License Agreement but failed to do so. Informatica's LMR was already able to calculate the number of physical cores on which Informatica Software was installed as evidenced by the "Cores" column in the 2016 LMR discussed above. In order to automatically enforce the License Agreement, Informatica would simply have to record the number of CPU cores allowed in the License Agreement and then use the functionality already demonstrated in the LMR to automatically check if the number of CPU cores on which the Informatica Software was deployed

[23] See, e.g., Deposition of Dwight Groff, January 17, 2018, 398:13 – 15.

exceeded the number of CPU cores allowed by the License Agreement. If the Informatica Software was deployed on more CPU cores that was permitted by the License Agreement, then such a system could take one or more automatic actions to enforce the License Agreement, including, for example:

i. Emailing and/or otherwise notifying one or more Informatica employees

ii. Emailing and/or otherwise notifying one or more RELX employees

iii. Setting a time period after which the Informatica Software would not execute on RELX machines in excess of the CPU cores allowed

iv. Immediately preventing the Informatica Software from executing on RELX machines in excess of the CPU cores allowed

47. This automatic approach to monitoring compliance with license agreements has been employed by other companies. As one example, Symantec's document entitled "The Symantec Approach to Software Asset Management" states that Symantec utilizes "Automation Policies" that can execute automated actions including "running a license compliance report" and "sending an email alert to the software administrator."[24] As another example, Veeam runs an automated process that "collects statistics on the current license usage and sends it periodically to the Veeam License Update server" and "runs in the background mode, once a week at a random time and day."[25] Informatica could have employed a similar automatic license monitoring system just as Symantec and Veeam have done.

48. Informatica employs an automatic license monitoring system in its Identity Resolution product but did not employ such a license monitoring solution for the Informatica

[24] See page 5 of https://www.symantec.com/content/dam/symantec/docs/solution-briefs/the-symantec-approach-to-software-asset-management-en.pdf, attached to this Report as Exhibit D.

[25] See https://helpcenter.veeam.com/docs/backup/cloud/sp_license_usage_report_online html?ver=95, attached to this Report as Exhibit E.

Software at issue in this case even though it easily could have done so. For example, in its online knowledgebase Informatica has a Frequently Asked Question that asks, "How is the CPU count for [Informatica Identity Resolution] licensing calculated?" Part of the answer that Informatica supplies to that question is, "The license server will use the second line (attribute cpu=<n>) to check for the number of allowed CPUs in a server. If this value exceeds, the IIR servers will not start and licensing errors will be seen in the ids.dbg file."[26] Informatica could have employed a similar automatic license monitoring system in the Informatica Software at issue in this case just as it had done for its Identity Resolution product.

VIII. RELX ACCRUED NO BENEFIT FROM THE INSTALLATION OF THE INFORMATICA SOFTWARE ON COMPUTERS HAVING A COLLECTIVE NUMBER OF CPU CORES GREATER THAN RELX HAD LICENSED

A. Informatica's Sizing Model for RELX Calculates at Most 12 CPU Cores as the Appropriate Number for RELX's Needs, Far Fewer than the Amount Informatica Sold to RELX and More than Double what RELX is Known to Have Utilized on Average during the Relevant Time Period

49. The spreadsheet with a filename of powercenter_sizing_model_LEXISNEXIS_V2.xls produced by Informatica as INFA_0000218583 ("Informatica's Sizing Model for RELX") appears to be a spreadsheet that was designed by Informatica to calculate the number of CPU cores and amount of memory (RAM) that computers would need in order to run the Informatica Software, and as such should have been taken into account by Informatica when determining the number of CPU core licenses that were appropriate for RELX's system that would run the Informatica Software. The calculations in Informatica's Sizing Model for RELX are based primarily on how much data is expected to be processed per hour and how many simultaneous data streams are intended to be

[26] See https://kb.informatica.com/faq/7/Pages/17/500713.aspx, attached to this Report as Exhibit F.

processed. In particular, Informatica's Sizing Model for RELX states on the spreadsheet tab labeled "Introduction":

> This sizing model attempts to isolate a few key factors that have the greatest bearing on the sizing estimate and can usually be ascertained with some certainty at the beginning of a data integration process. By focusing on these key factors and using a 'best guess' of their value or expected values - a reasonable high level hardware recommendation will be generated.

Informatica's Sizing Model for RELX also states on the spreadsheet tab labeled "Calculations Explained": "The key driving factors for calculating the base CPU core size are 'core MB per sec' (data rate) and 'core per session' (job load)."

50. Informatica's Sizing Model for RELX calculates the number of CPU cores required by taking client-specific inputs (e.g. from RELX) from the spreadsheet tab entitled "Sizing Parameters" (depicted below in Fig. 6, Fig. 8, and Fig. 10), combining them with internal parameters recorded in the spreadsheet tab entitled "Internal Sizing Params" and the Factors and Description portion of the spreadsheet tab entitled "Sizing Calculation," and then outputting the calculation in the "Informatica PowerCenter Sizing Results" portion of the spreadsheet tab entitled "Sizing Calculations" (depicted in Fig. 7, Fig. 9, and Fig. 11). The primary client-specific inputs are the "number of gigabytes per hour" and the "number of simultaneous jobs on average." These values are inputted using the fields under "Method 1" of the "Sizing Parameters" spreadsheet tab or they themselves are calculated using the values inputted using the fields under "Method 2" or "Method 3" of the "Sizing Parameters" spreadsheet tab. These primary client-specific inputs are used to calculate the "Initial CPU" value in the "Sizing Calculation" spreadsheet tab.

51. The "Initial CPU" value is then automatically modified by several secondary client-specific inputs, which include (1) Processor Type, (2) Continuous and/or Real Time, (3)

Load Window Criticality, (4) Aggregates and Sorting, (5) Data Volume Growth, (6) ASCII or Unicode Data, (7) Unix or Windows, (8) Data Transformation (B2B DT), and (9) Lookup Sizing. Each of the values entered for the secondary client-specific inputs are converted to a percentage, added together, and then multiplied by the "Initial CPU" value to adjust the calculation, resulting in the "Adjusted CPU" value.[27] The "Adjusted CPU" value is then rounded to an even number to determine the low end of the range of the final calculation for the number of CPU cores, and the high end of the range of the final calculation for the number of CPU cores is determined to be two more than the low end of the range.

52. Using several different methodologies to generate reasonable input parameters for Informatica's Sizing Model for RELX, Informatica's Sizing Model for RELX calculates at most 12 CPU cores as the appropriate number for RELX's needs, far fewer than the amount Informatica sold to RELX (72 during certain time periods) and more than double the CPU Core Equivalents (defined below) that RELX needed on average (between 5.1 and 5.5) to process files on the ICCE Platform over the Relevant Time Period for which utilization data is available. Informatica did not adhere to the calculations of CPU cores determined in Informatica's Sizing Model for RELX, instead selling RELX far more CPU core licenses than Informatica's Sizing Model for RELX calculates using the parameters below. A summary of the results of the three approaches used in this Report to calculate CPU cores using Informatica's Sizing Model for RELX is below.

[27] The exception is the "Processor Type" parameter, which is used to scale the "Initial CPU" value before the other secondary client-specific inputs are taken into account.

Approach	Maximum Number of CPU Cores Calculated by Informatica's Sizing Model for RELX
Use the parameters in Informatica's Sizing Model for RELX as produced.	6
Use the Maximal Parameters (defined below) specified by Mr. Groff.	10
Use the Maximal Parameters and estimates specified by Mr. Groff in conjunction with historical file processing data from the ICCE Files Processed Report.	12

1. As Produced, Informatica's Sizing Model for RELX Calculates a Maximum of Six CPU Cores as the Appropriate Number of CPU Cores for RELX's Needs

53. Informatica's Sizing Model for RELX appears to have been produced with parameters for calculating the number of CPU cores that RELX would need to run the Informatica Software effectively. For example, the filename of Informatica's Sizing Model for RELX, "powercenter_sizing_model_LEXISNEXIS_V2.xls," includes "LEXISNEXIS." Based on the parameters in the document as produced, Informatica's Sizing Model for RELX calculates between four and six CPU cores for RELX's data processing needs in order to process 18 gigabytes of data per hour across seven simultaneous jobs as depicted in Fig. 6 and Fig. 7 below.

Informatica PowerCenter Sizing Questions

Data Volume Rate

Data volumes are a critical aspect as the CPU cycles must be available to handle the data volume in the appropriate timeframe. Getting a reasonable estimate of the volume of data to be moved on a nightly/daily basis is the cornerstone of a sizing effort.

Method 1 (Volume based)	
Number of Gigabytes per hour	18
Number of simultaneous jobs on average?	7
Method 2 (Existing load process)	
How many loads?	
How much data is being moved (in GB)?	
What is your load window in minutes?	
Method 3 (Expected load process)	
How many target tables do you have to load?	
Size of data to load (source data) in GB?	
What is the load window in minutes?	

Processor Type

Chose the processor. The "SPECint2006 Rate" values from 2009-2011 benchmarks are used internally

What processor do you plan to use	IBM Power7 (3.61 GHz)

Continuous and/or Real Time

The assumption is that continuous and/or real-time workloads will require more CPU and memory. This is because there is less flexibility in workload management. RT sessions must run, and they must run now and they should not be slowed by other processes. On the other hand, RT sessions often have much lower throughput requirements since they are always on (hence, large load window) so the estimate will be adjusted accordingly.

What percent of sessions/loads will be real time?	

Load Window Criticality

How critical is it that the load windows is always met?	Somewhat

Aggregates and Sorting

Aggregates and sorters require additon CPU cores and RAM

What is the expected use of aggregates/sorters?	Low

Data Volume Growth

What is the expected yearly data growth (%)?	25%

ASCII or Unicode Data

Percentage of data expected to be Unicode	Low

Unix or Windows

Are you planning to use Windows or Unix/Linux	Unix

Data Transformation (B2B DT)

What percentage of sessions/loads will use B2B DT?	Low

Lookup Sizing

Lookups (caching data tables to match values) require additional CPU cores and RAM. It is an important factor in sizing the box. Use an educated guess as to the size of your lookup requirements. If you are loading a warehouse, think in terms of the size of the dimension tables you might cache when loading a fact table for example.

Percent of lookups with > 500K rows?	Medium

Application Type(s)

What sort of application(s) will PowerCenter be used for?

Other Considerations

Please include any other considerations you feel are important to the sizing effort. Any environmental information, restrictions, needs should be listed here.

Fig. 6: Informatica's Sizing Model for RELX with Parameters as Produced

Informatica PowerCenter Sizing Results

Component	Details	Ram Factor	Initial CPU	Adjusted CPU
Data Volume Rate Method 1	Using CPU/MB sec factor	9.2	0.5	2.3
Data Volume Rate Method 2	Using CPU/MB sec factor	0	0	0
Data Volume Rate Method 3	Using CPU/MB sec factor	0	0	0
Base Size		**9.2**	**0.5**	**2.3**
Continuous And/Or Real Time	Ranges (0,25%,60%,100%)	0%	0%	0%
Load Window Criticality	Ranges (-30%,0,50%,75%)	0%	0%	0%
Aggregates and Sorting	Ranges(-20%, 0, 50%)	-20%	-20%	-20%
Data Volume Growth		50%	50%	50%
ASCII or Unicode Data	Ranges(0,25%,60%,100%)	30%	30%	30%
Operating System	Unix/Linux or Windows	0%	0%	0%
B2B DT	Ranges(0, 50%, 100%)	0%	0%	0%
Lookup Sizing	Ranges (Ram = -20%,0,50%)	0%	0%	0%
Application Types + Other	Subjective Factor (in %)			
Total Adjustment Factor		**60%**	**60%**	**60%**
Final Sizing Raw		**14.72**	**0.8**	**3.68**
Final Sizing Adjusted		**16**	**2**	**4**
Sizing Upper Range		**24**	**4**	**6**
Minimum Number of Nodes				**1**

Informatica recommends at least 1 server(s) with a total of 4 to 6 CPU cores and 16 to 24 GB of RAM

03/2012 Note: The Orace/Sun T-Series hardware does not scale well with PowerCenter. Please use M-Series class hardware.

Factors	Description
5	Job Length Average
0	Est Concurrent Jobs using JLA
0	Est Concurrent Jobs using JLA
4	RAM per core Factor
5.5	Core MB per second
0.5	Core per Session Sizing Factor
100	IO bus clock speed (MHz)
2.12	Core Adjust Factor (Do not edit)
0.25	Windows vs Unix CPU core Factor
196	Max Windows Useable Memory (GB)
2	Data Volume Growth Time (yrs)
2	Max Adjust CPU Factor

Fig. 7: Informatica's Sizing Model for RELX Results with Parameters as Produced

2. Based on the Maximal Parameters Specified by Mr. Groff, Informatica's Sizing Model for RELX Calculates a Maximum of Ten CPU Cores as the Appropriate Number of CPU Cores for RELX's Needs

54. In order to determine if Informatica's Sizing Model for RELX might have reasonably predicted a different number of CPU cores for RELX to use with Informatica's Software, I asked Dwight Groff, an employee of RELX who is knowledgeable about RELX's implementation of the Informatica Software, to provide parameters for the Informatica's Sizing Model for RELX. I provided Mr. Groff with only the "Sizing Parameters" tab of Informatica's Sizing Model for RELX, which included instructions for how to fill in the parameters. I also included on the Sizing Parameters tab the field value choices present in Informatica's Sizing Model for RELX's "Internal Sizing Params" tab so that Mr. Groff could select among these provided choices for each parameter. The spreadsheet I provided to Mr. Groff is attached to this Report as Exhibit G. Mr. Groff provided his response to me via counsel in an email ("Groff Email"), which is attached to this Report as Exhibit H and includes the prose of my initial email to Mr. Groff, his response, as well as the contents of the spreadsheet Mr. Groff included in his response. Mr. Groff provided parameters for Informatica's Sizing Model for RELX across three different time spans, and across those different time spans I selected the parameters ("Maximal Parameters") for each field that would result in Informatica's Sizing Model for RELX calculating the largest number of CPU cores. Based on the Maximal Parameters in the Groff Email that Mr. Groff provided, Informatica's Sizing Model for RELX calculates between eight and ten CPU cores for RELX's data processing needs in order to process 0.8 gigabytes of data per hour across 30 simultaneous jobs as depicted in Fig. 8 and Fig. 9 below.

Informatica PowerCenter Sizing Questions

Data Volume Rate

Data volumes are a critical aspect as the CPU cycles must be available to handle the data volume in the appropriate timeframe. Getting a reasonable estimate of the volume of data to be moved on a nightly/daily basis is the cornerstone of a sizing effort.

Method 1 (Volume based)	
Number of Gigabytes per hour	0.8
Number of simultaneous jobs on average?	30
Method 2 (Existing load process)	
How many loads?	
How much data is being moved (in GB)?	
What is your load window in minutes?	
Method 3 (Expected load process)	
How many target tables do you have to load?	
Size of data to load (source data) in GB?	
What is the load window in minutes?	

Processor Type

Chose the processor. The "SPECint2006 Rate" values from 2009-2011 benchmarks are used internally

What processor do you plan to use	Intel Xeon (2.4 GHz)

Continuous and/or Real Time

The assumption is that continuous and/or real-time workloads will require more CPU and memory. This is because there is less flexibility in workload management. RT sessions must run, and they must run now and they should not be slowed by other processes. On the other hand, RT sessions often have much lower throughput requirements since they are always on (hence, large load window) so the estimate will be adjusted accordingly.

What percent of sessions/loads will be real time?	High

Load Window Criticality

How critical is it that the load windows is always met?	Somewhat

Aggregates and Sorting

Aggregates and sorters require additon CPU cores and RAM

What is the expected use of aggregates/sorters?	Medium

Data Volume Growth

What is the expected yearly data growth (%)?	50%

ASCII or Unicode Data

Percentage of data expected to be Unicode	None

Unix or Windows

Are you planning to use Windows or Unix/Linux	Unix

Data Transformation (B2B DT)

What percentage of sessions/loads will use B2B DT?	High

Lookup Sizing

Lookups (caching data tables to match values) require additional CPU cores and RAM. It is an important factor in sizing the box. Use an educated guess as to the size of your lookup requirements. If you are loading a warehouse, think in terms of the size of the dimension tables you might cache when loading a fact table for example.

Percent of lookups with > 500K rows?	Low

Application Type(s)

What sort of application(s) will PowerCenter be used for?

Other Considerations

Please include any other considerations you feel are important to the sizing effort. Any environmental information, restrictions, needs should be listed here.

Fig. 8: Informatica's Sizing Model for RELX with Maximal Parameters from Mr. Groff

Informatica PowerCenter Sizing Results

Component	Details	Ram Factor	Initial CPU	Adjusted CPU
Data Volume Rate Method 1	Using CPU/MB sec factor	8.4	0.1	2.1
Data Volume Rate Method 2	Using CPU/MB sec factor	0	0	0
Data Volume Rate Method 3	Using CPU/MB sec factor	0	0	0
Base Size		**8.4**	**0.1**	**2.1**
Continuous And/Or Real Time	Ranges (0,25%,60%,100%)	100%	100%	100%
Load Window Criticality	Ranges (-30%,0,50%,75%)	0%	0%	0%
Aggregates and Sorting	Ranges(-20%, 0, 50%)	0%	0%	0%
Data Volume Growth		100%	100%	100%
ASCII or Unicode Data	Ranges(0,25%,60%,100%)	0%	0%	0%
Operating System	Unix/Linux or Windows	0%	0%	0%
B2B DT	Ranges(0, 50%, 100%)	0%	100%	100%
Lookup Sizing	Ranges (Ram = -20%,0,50%)	-20%	-10%	-10%
Application Types + Other	Subjective Factor (in %)			
Total Adjustment Factor		**180%**	**290%**	**290%**
Final Sizing Raw		**23.52**	**0.39**	**8.19**
Final Sizing Adjusted		**24**	**0**	**8**
Sizing Upper Range		**30**	**2**	**10**
Minimum Number of Nodes				**1**

Informatica recommends at least 1 server(s) with a total of 8 to 10 CPU cores and 24 to 30 GB of RAM

03/2012 Note: The Orace/Sun T-Series hardware does not scale well with PowerCenter. Please use M-Series class hardware.

Factors	Description
5	Job Length Average
0	Est Concurrent Jobs using JLA
0	Est Concurrent Jobs using JLA
4	RAM per core Factor
5.5	Core MB per second
0.5	Core per Session Sizing Factor
100	IO bus clock speed (MHz)
1.58	Core Adjust Factor (Do not edit)
0.25	Windows vs Unix CPU core Factor
196	Max Windows Useable Memory (GB)
2	Data Volume Growth Time (yrs)
2	Max Adjust CPU Factor

Fig. 9: Informatica's Sizing Model for RELX Results with Maximal Parameters from Mr. Groff

3. Based on an Estimated File Size and Other Parameters Specified by Mr. Groff as well as Historical File Totals Processed by the Informatica Software, Informatica's Sizing Model for RELX Calculates a Maximum of 12 CPU Cores as the Appropriate Number of CPU Cores for RELX's Needs

55. In order to further determine if Informatica's Sizing Model for RELX might have reasonably predicted a different number of CPU cores for RELX to use with Informatica's Software, I asked Mr. Groff to provide an estimate of the average file size for the files that the Informatica Software processed while in operation on RELX systems. Mr. Groff's estimate was 150 kilobytes as indicated in the Groff Email. To determine the total number of files processed, I referred to the ICCE Files Processed Report, which includes daily totals of files processed by the ICCE Platform, which included the Informatica Software. In total 974,179,899 files were processed over 1,827 days. On average that's about 79,981,929 kilobytes of data processed per day,[28] or about 3.2 gigabytes per hour.[29] I then used this modified estimate of data throughput per hour along with the Maximal Parameters provided by Mr. Groff to determine a new calculation for the number of CPU cores using Informatica's Sizing Model for RELX. Based on the average file size and the Maximal Parameters that Mr. Groff provided, Informatica's Sizing Model for RELX calculates between 10 and 12 CPU cores for RELX's data processing needs in order to process 3.2 gigabytes of data per hour across 30 simultaneous jobs as depicted in Fig. 10 and Fig. 11 below.

[28] 150 kilobytes per file * 974,179,899 files / 1,827 days

[29] 79,981,929 kilobytes per day / 1024 kilobytes per megabyte / 1024 megabytes per gigabyte / 24 hours per day

Informatica PowerCenter Sizing Questions

Data Volume Rate

Data volumes are a critical aspect as the CPU cycles must be available to handle the data volume in the appropriate timeframe. Getting a reasonable estimate of the volume of data to be moved on a nightly/daily basis is the cornerstone of a sizing effort.

Method 1 (Volume based)	
Number of Gigabytes per hour	3.2
Number of simultaneous jobs on average?	30
Method 2 (Existing load process)	
How many loads?	
How much data is being moved (in GB)?	
What is your load window in minutes?	
Method 3 (Expected load process)	
How many target tables do you have to load?	
Size of data to load (source data) in GB?	
What is the load window in minutes?	

Processor Type

Chose the processor. The "SPECint2006 Rate" values from 2009-2011 benchmarks are used internally

What processor do you plan to use	Intel Xeon (2.4 GHz)

Continuous and/or Real Time

The assumption is that continuous and/or real-time workloads will require more CPU and memory. This is because there is less flexibility in workload management. RT sessions must run, and they must run now and they should not be slowed by other processes. On the other hand, RT sessions often have much lower throughput requirements since they are always on (hence, large load window) so the estimate will be adjusted accordingly.

What percent of sessions/loads will be real time?	High

Load Window Criticality

How critical is it that the load windows is always met?	Somewhat

Aggregates and Sorting

Aggregates and sorters require additon CPU cores and RAM

What is the expected use of aggregates/sorters?	Medium

Data Volume Growth

What is the expected yearly data growth (%)?	50%

ASCII or Unicode Data

Percentage of data expected to be Unicode	None

Unix or Windows

Are you planning to use Windows or Unix/Linux	Unix

Data Transformation (B2B DT)

What percentage of sessions/loads will use B2B DT?	High

Lookup Sizing

Lookups (caching data tables to match values) require additional CPU cores and RAM. It is an important factor in sizing the box. Use an educated guess as to the size of your lookup requirements. If you are loading a warehouse, think in terms of the size of the dimension tables you might cache when loading a fact table for example.

Percent of lookups with > 500K rows?	Low

Application Type(s)

What sort of application(s) will PowerCenter be used for?

Other Considerations

Please include any other considerations you feel are important to the sizing effort. Any environmental information, restrictions, needs should be listed here.

Fig. 10: Informatica's Sizing Model for RELX with Maximal Parameters from Mr. Groff with Updated GB/hour

Informatica PowerCenter Sizing Results

Component	Details	Ram Factor	Initial CPU	Adjusted CPU
Data Volume Rate Method 1	Using CPU/MB sec factor	8.8	0.2	2.2
Data Volume Rate Method 2	Using CPU/MB sec factor	0	0	0
Data Volume Rate Method 3	Using CPU/MB sec factor	0	0	0
Base Size		**8.8**	**0.2**	**2.2**
Continuous And/Or Real Time	Ranges (0,25%,60%,100%)	100%	100%	100%
Load Window Criticality	Ranges (-30%,0,50%,75%)	0%	0%	0%
Aggregates and Sorting	Ranges(-20%, 0, 50%)	0%	0%	0%
Data Volume Growth		100%	100%	100%
ASCII or Unicode Data	Ranges(0,25%,60%,100%)	0%	0%	0%
Operating System	Unix/Linux or Windows	0%	0%	0%
B2B DT	Ranges(0, 50%, 100%)	0%	100%	100%
Lookup Sizing	Ranges (Ram = -20%,0,50%)	-20%	-10%	-10%
Application Types + Other	Subjective Factor (in %)			
Total Adjustment Factor		**180%**	**290%**	**290%**
Final Sizing Raw		**24.64**	**0.78**	**8.58**
Final Sizing Adjusted		**26**	**2**	**10**
Sizing Upper Range		**32**	**4**	**12**
Minimum Number of Nodes				**1**

Informatica recommends at least 1 server(s) with a total of 10 to 12 CPU cores and 26 to 32 GB of RAM

03/2012 Note: The Orace/Sun T-Series hardware does not scale well with PowerCenter. Please use M-Series class hardware.

Factors	Description
5	Job Length Average
0	Est Concurrent Jobs using JLA
0	Est Concurrent Jobs using JLA
4	RAM per core Factor
5.5	Core MB per second
0.5	Core per Session Sizing Factor
100	IO bus clock speed (MHz)
1.58	Core Adjust Factor (Do not edit)
0.25	Windows vs Unix CPU core Factor
196	Max Windows Useable Memory (GB)
2	Data Volume Growth Time (yrs)
2	Max Adjust CPU Factor

Fig. 11: Informatica's Sizing Model for RELX Results with Maximal Parameters from Mr. Groff with Updated GB/hour

B. The ICCE Server Utilization Report and the ICCE Files Processed Report Confirm that RELX Never Needed more than 56 CPU Core Licenses to Meet its Needs During the Relevant Time Period

1. Overview of the ICCE Server Utilization Report

56. The ICCE Server Utilization Report is a spreadsheet that shows historical CPU utilization for the seven servers ("ICCE Servers") that RELX used at various times to run its ICCE Platform, which used Informatica Software for certain components. It shows CPU utilization over the time period of December 9, 2014 to November 7, 2017, which is approximately the last date on which RELX processed any files with the ICCE Platform.[30] My understanding from counsel is that the ICCE Server Utilization Report does not contain data prior to December 9, 2014 because the data used to generate it was periodically automatically overwritten, and when the data used to generate it was first gathered for this litigation, December 9, 2014 was the earliest date for which such data still existed. The ICCE Server Utilization Report shows for each of the seven ICCE Servers the average CPU utilization across all of the given ICCE Server's cores for each hour[31] of each day, as reported by the ICCE Server's operating system. It does not show the utilization of any individual cores and does not show average utilization at any temporal granularity smaller than one hour. My understanding from discussions with Dwight Groff[32] is that the RELX did not accrue a benefit from processing any document even three hours more quickly, and therefore the hour-by-hour granularity of the ICCE Server Utilization Report is a sufficient granularity at which to examine the extent to which RELX accrued a benefit. Finally, the utilization reported in the ICCE Server Utilization Report includes

[30] The last date of core utilization from the ICCE Server Utilization Report is November 7, 2017, and the last date that a file was processed as part of the ICCE Platform according to the ICCE Files Processed Report is November 8, 2017.

[31] Across all seven ICCE Servers, more than 99% of hours have data reported over the approximately three-year timeframe covered by the ICCE Server Utilization Report.

[32] See, e.g., the Groff Email.

time that cores were executing any part of the ICCE Platform software, including RELX custom software, in the ICCE Platform, not only the Informatica Software.

57. "Utilization" in the ICCE Server Utilization Report refers to a percentage-based metric used to measure over a given time period the number of instruction-executing cycles[33] a core actually performs as a percentage of the maximum number of instruction-executing cycles a core is rated to perform over that same time period. The rating of a given core is also commonly referred to as its speed and modern cores usually have speeds measured in gigahertz (cycles per second). As an example, if a given core has a speed of 1 gigahertz (1 billion cycles per second), and over the course of an hour (comprising 3,600 seconds), it performs 3,600 billion instruction-executing cycles, then it was 100% utilized on average for that hour. If that same core performed 1,800 billion instruction-executing cycles over an hour, then it was 50% utilized on average for that hour.

58. Multiple cores on the same CPU can work together to execute all the instructions given to the CPU by computer programs even if the work could be also done in the same amount of time by fewer cores. Continuing the example above, two cores both with speeds of 1 gigahertz could divide up the execution of instructions requiring 3,600 billion cycles over the course of an hour in various ways. If they both perform 1,800 billion instruction-executing cycles, then both cores individually were 50% utilized on average over the hour as well as collectively 50% utilized on average over the course of the hour. Depending on the circumstances and particular configurations, however, the CPU might divide the instructions between the two cores differently. One core might even perform all 3,600 billion instruction-executing cycles to execute the instructions over the course of the hour, while the other performed zero instruction-executing

[33] Depending on the CPU architecture, instructions can take one or more cycles to execute. When a core does not have any instructions to execute from computer programs, it can instead perform instructions for an "idle task" that does not contribute to the execution of any computer program being executed by the core.

cycles, effectively using only one core. In that case one core would be 100% utilized and the other would be 0% utilized over the hour, but collectively they would still be 50% utilized. Another way to express this is to say that the 3,600 billion instruction-executing cycles performed over that hour required the *equivalent* of one core rated for 1 gigahertz operating at 100% utilization over that hour, even if the instruction-executing cycles may have actually been distributed over multiple cores. In this Report, I use the term "CPU Core Equivalents" to refer to this measurement of how many cores operating at 100% utilization it would take to perform a given number of instruction-executing cycles over a given time period.

59. CPUs can have multiple cores, computers can have multiple CPUs, and distributed computing systems can have multiple computers (also known as "servers") configured to execute the instructions of various computer programs. The ICCE Platform was configured as a distributed computing system using the ICCE Servers, each with their own configuration of cores. The various programs executed on the ICCE Platform might at any time execute on any one of the ICCE Servers, on any one of that ICCE Server's CPUs, and on any one of that CPU's cores, depending on the circumstances and configurations of that CPU, its ICCE Server, and the ICCE Platform as a whole.

2. Historical Utilization Data Confirms that RELX Never Needed more than 56 CPU Core Licenses to Meet its Needs During the Late Portion of the Relevant Time Period During Which RELX Had 56 Cores Deployed

60. During the time period ("Late Time Period") of the 8th hour of May 15, 2016 through November 7, 2017, approximately the last date on which RELX processed any files with the ICCE Platform, the ICCE Platform had four ICCE Servers configured. The ICCE Servers psc3813, psc3814, and psc33817 each had two CPUs with four cores per CPU, for a total of eight

cores per server.[34] The ICCE Server psc33841had four CPUs with eight cores per CPU, for a total of 32 cores.[35] For the entirety of the Late Time Period, the ICCE Platform was configured with the 56 cores from these servers.[36] During the Late Time Period RELX had licenses for 56 CPU cores.[37]

61. According to the data in the ICCE Server Utilization Report, during the Late Time Period, the hourly weighted average[38] core utilization was approximately 9.3%, and the hourly weighted average CPU Core Equivalents was approximately 5.1.[39] The single hour during which the hourly weighted average core utilization and hourly weighted average CPU Core Equivalents was highest over this time period is the 16th hour of June 7, 2017, which had an hourly weighted average core utilization of approximately 96.3% and an hourly weighted average CPU Core Equivalents of approximately 53.9, which is less than the 56 CPU core licenses that RELX possessed during this time period.

62. Below are graphs that show the hourly weighted average core utilization and hourly weighted average CPU Core Equivalents for the Late Time Period.

[34] See, e.g., Deposition of Dwight Groff, January 17, 2018, 346:5 – 347:9 and 492:3 – 5, and the Dell R710 Spec Sheet at https://www.dell.com/downloads/global/products/pedge/r710-spec-sheet.pdf and attached to this Report as Exhibit I.

[35] See, e.g., Deposition of Dwight Groff, January 17, 2018, 492:6 – 8 and the Dell R910 Spec Sheet at https://www.dell.com/downloads/global/products/pedge/r910-spec-sheet.pdf and attached to this Report as Exhibit J.

[36] 8 + 8 + 8 + 32

[37] See, e.g., Exhibit K to the Reihl Declaration.

[38] The average is weighted to account for the fact that some ICCE Servers had more cores than others and the data in the ICCE Server Utilization Report is reported on a per-server and not a per-CPU or per-core basis.

[39] ~5.1 CPU Core Equivalents = ~9.3% of 56 cores available.
Generally for this time period, CPU Core Equivalents = Weighted Average Core Utilization * 56 / 100.

CONFIDENTIAL

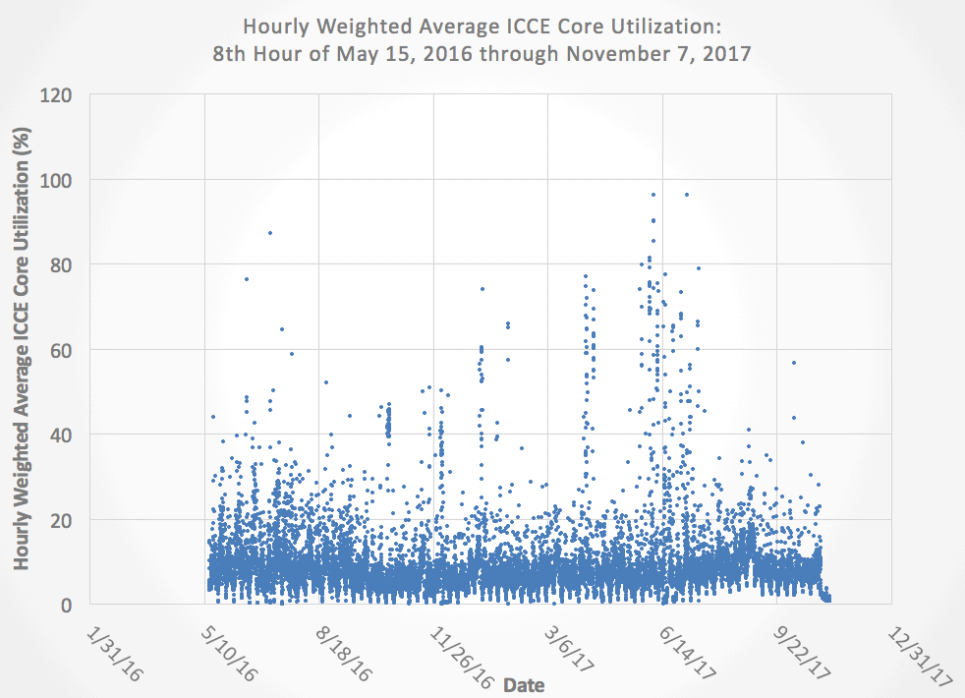

Fig. 12

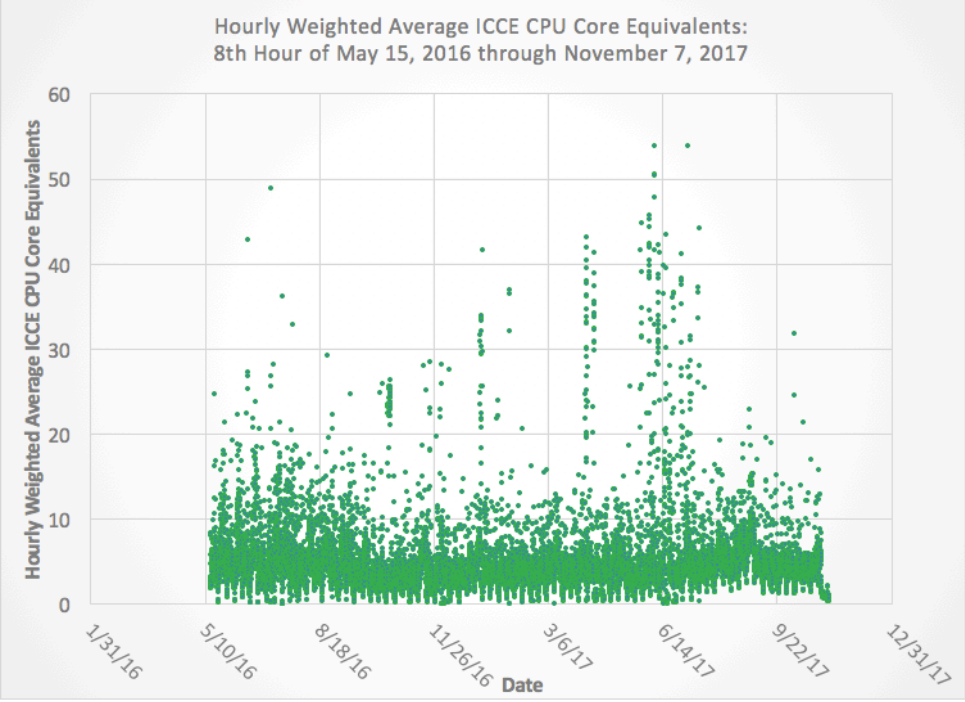

Fig. 13

63. During the Late Time Period, RELX was licensed for 56 CPU cores, and RELX's ICCE Platform was configured with 56 cores. During this same time period, the hourly weighted average CPU Core Equivalents was approximately 5.1 with an hourly average maximum of 53.9.

3. Historical Utilization Data Confirms that RELX Never Needed more than 56 CPU Core Licenses to Meet its Needs During the Middle Portion of the Relevant Time Period During Which RELX Had 104 Cores Deployed

64. During the time period ("Middle Time Period") of December 9, 2014 (the earliest date for which there is data in the ICCE Server Utilization Report) through the 7th hour of May 15, 2016, the ICCE Platform had seven ICCE Servers configured. The ICCE Servers psc3813, psc3814, psc3815, psc3816, and psc33817 each had two CPUs with four cores per CPU, for a total of eight cores per server.[40] The ICCE Servers psc33841 and psc33842 each had four CPUs with eight cores per CPU, for a total of 32 cores per server.[41] For the duration of the Middle Time Period, the ICCE Platform was configured with 104 cores from these servers.[42] From December 9, 2014 through February 10, 2015 RELX had licenses for 72 CPU cores,[43] and from February 11, 2015 through May 15, 2016 RELX had licenses for 56 CPU cores.[44]

65. According to the data in the ICCE Server Utilization Report, during the Middle Time Period, the hourly weighted average core utilization was approximately 5.3%, and the

[40] See, e.g., Deposition of Dwight Groff, January 17, 2018, 346:5 – 347:9 and 492:3 – 5, and the Dell R710 Spec Sheet at https://www.dell.com/downloads/global/products/pedge/r710-spec-sheet.pdf and attached to this Report as Exhibit I.

[41] See, e.g., Deposition of Dwight Groff, January 17, 2018, 492:6 – 8 and the Dell R910 Spec Sheet at https://www.dell.com/downloads/global/products/pedge/r910-spec-sheet.pdf and attached to this Report as Exhibit J.

[42] 8 + 8 + 8 + 8 + 8 + 32 + 32

[43] See, e.g. Exhibit E to the Reihl Declaration.

[44] See, e.g. Exhibit K to the Reihl Declaration.

hourly weighted average CPU Core Equivalents was approximately 5.5.[45] The single hour during which the hourly weighted average core utilization and hourly weighted average CPU Core Equivalents was highest over the Middle Time Period is the 7th hour of March 17, 2016, which had an hourly weighted average core utilization of approximately 51.6% and an hourly weighted average CPU Core Equivalents of approximately 53.6, which is less than the minimum 56 CPU core licenses that RELX possessed during most of this time period.

66. Below are graphs that show the hourly weighted average core utilization and hourly weighted average CPU Core Equivalents for this time period.

[45] ~5.5 CPU Core Equivalents = ~5.3% of 104 cores available.
Generally for this time period, CPU Core Equivalents = Weighted Average Core Utilization * 104 / 100.

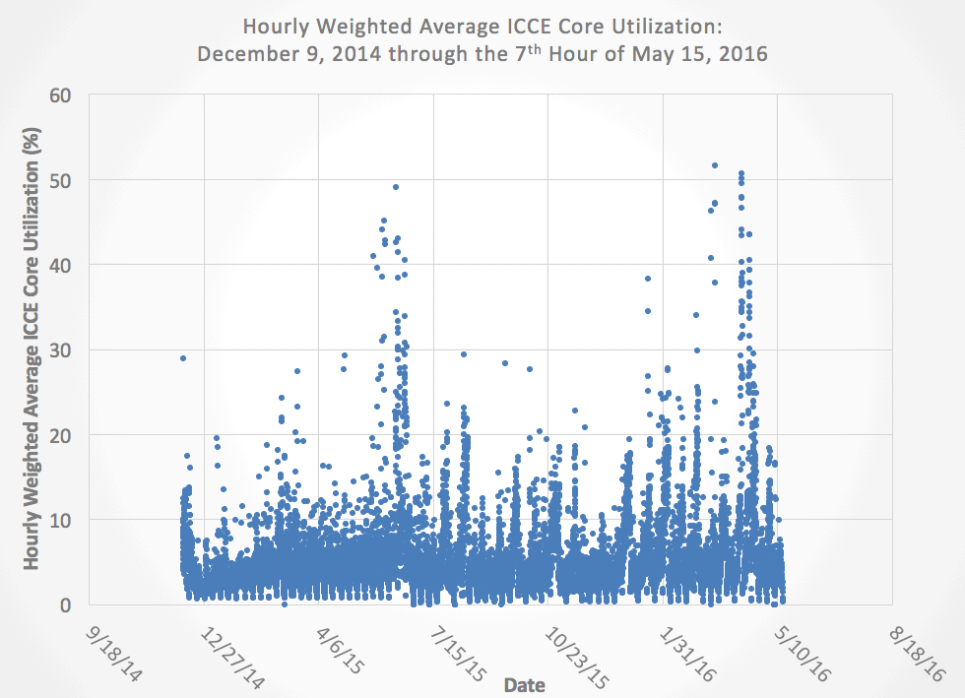
Hourly Weighted Average ICCE Core Utilization:
December 9, 2014 through the 7th Hour of May 15, 2016
Hourly Weighted Average ICCE Core Utilization (%)
60
50
40
30
20
10
0
9/18/14
12/27/14
4/6/15
7/15/15
10/23/15
1/31/16
5/10/16
8/18/16
Date

Fig. 14

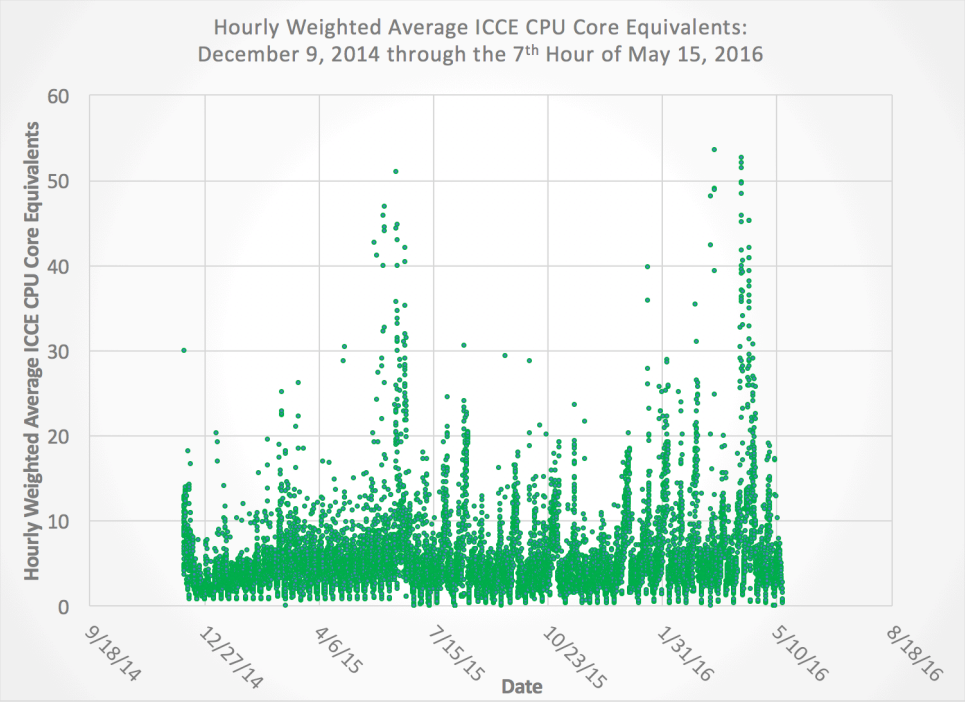
Hourly Weighted Average ICCE CPU Core Equivalents:
December 9, 2014 through the 7th Hour of May 15, 2016
Hourly Weighted Average ICCE CPU Core Equivalents
60
50
40
30
20
10
0
9/18/14
12/27/14
4/6/15
7/15/15
10/23/15
1/31/16
5/10/16
8/18/16
Date

Fig. 15

67. During the Middle Time Period, RELX accrued no benefit from the installation of the Informatica software on computers having a collective number of CPU cores greater than RELX had licensed. Over the course of every hour accounted for in the ICCE Server Utilization Report during the Middle Time Period, the CPU Core Equivalents was no more than 53.6, which is less than the minimum 56 CPU core licenses that RELX possessed during the Middle Time Period. Furthermore, the CPU Core Equivalents during the Middle Time Period are very similar to the CPU Core Equivalents used during the Late Time Period, discussed in the previous section. Over the Late Time Period the hourly weighted average CPU Core Equivalents was approximately 5.1, which is very close to the approximately 5.5 CPU Core Equivalents used during the Middle Time Period. Over the Late Time Period the maximum hourly weighted average CPU Core Equivalents was 53.9, which is very close to the approximately 53.6 maximum hourly weighted average CPU Core Equivalents for the Middle Time Period. If RELX's ICCE Platform had been configured with 56 cores during the Middle Time Period, it is very likely that its CPU Core Equivalents would have changed at most by a trivial amount, as evidenced by that very scenario occurring during the Late Time Period with a decrease in hourly weighted average CPU Core Equivalents from 5.5 to 5.1 even as the number of files processed by the ICCE Platform increased over the Middle Time Period and the Late Time Period, as described below.[46]

[46] The slight decrease in the hourly weighted average CPU Core Equivalents from the Middle Time Period to the Late Time Period despite the somewhat substantial increase in files processed by the ICCE Platform over the same time periods could be explained by many external factors including but not limited to the fact that RELX was making the ICCE Platform more efficient over time and decreasing the number of components that relied on Informatica's Data Transformation product.

4. Historical Utilization Data and Historical File Processing Data Indicate that RELX Never Needed more than 72 CPU Core Licenses During the Early Portion of the Relevant Time Period During Which RELX Had 104 or Fewer Cores Deployed

68. The ICCE Server Utilization Report does not contain data for the time period ("Early Time Period") of November 8, 2012 through December 8, 2014, and during the Early Time Period RELX had licenses for 72 CPU cores.[47] The ICCE Files Processed Report is a spreadsheet that shows the total number of files processed by the ICCE Platform, which used the Informatica Software for certain components, for each day from November 8, 2012 to November 7, 2017 (the Relevant Time Period), spanning the entirety of the Early Time Period, the Middle Time Period, and the Late Time Period. The daily average number of files processed by the ICCE Platform in the Early Time Period is 66,771, the daily average number of files processed by the ICCE Platform in the Middle Time Period is 577,073, and the daily average number of files processed by the ICCE Platform in the Late Time Period is 1,145,844. Since the instructions executed by the cores in the ICCE Platform were generally in service of processing files, and the average number of files processed per day during the Early Time Period is so much lower than the average number of files processed per day during the Middle Time Period and Late Time Period, it is very unlikely that core utilization during the Early Time Period was any greater than that of the Middle Time Period and the Later Time Period, and the hourly weighted average CPU Core Equivalents during the Early Time Period was therefore almost certainly not greater than either the hourly weighted average CPU Core Equivalents during the Middle Time Period or the Late Time Period.

69. Below is a graph that shows the daily total files processed by the ICCE Platform as well as average daily total files processed by the ICCE Platform for the Early Time Period, the

[47] See, e.g. Exhibit E to the Reihl Declaration.

Middle Time Period, and the Late Time Period. This graph also includes a linear trend line that shows that the daily total files processed by the ICCE Platform generally increased over the Relevant Time Period.

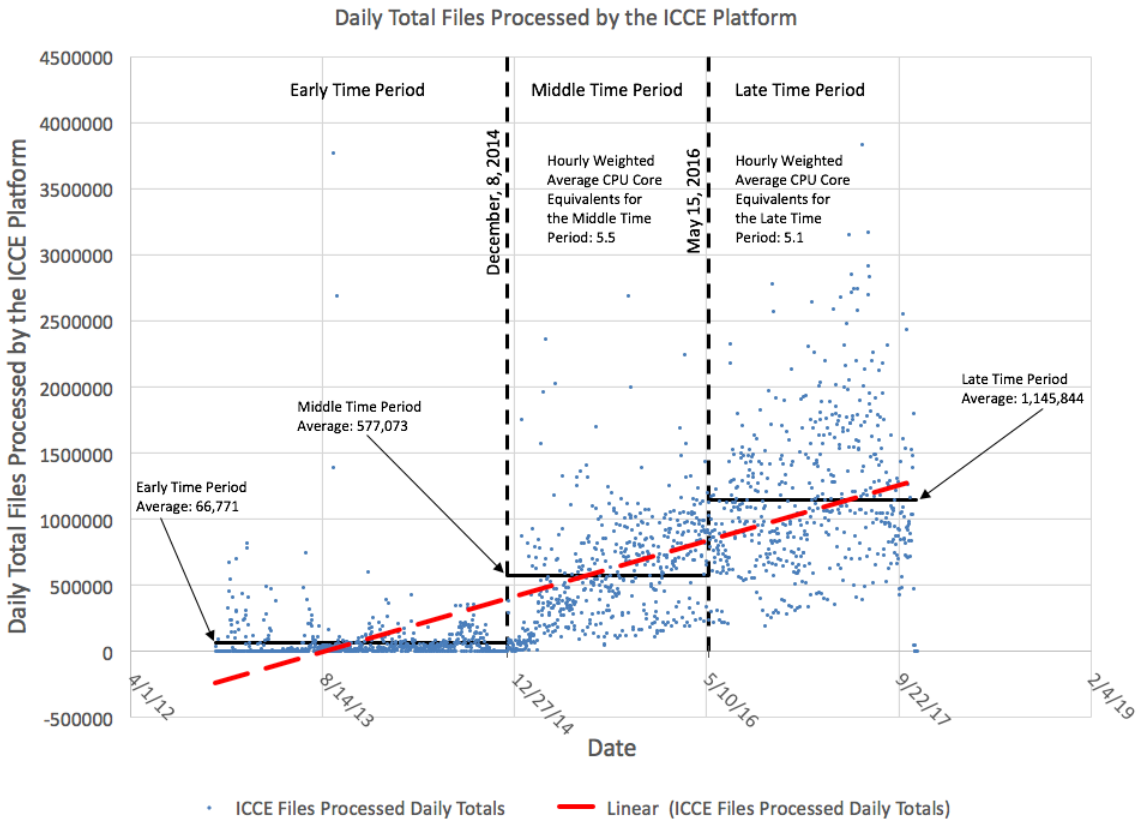

Fig. 16

IX. INFORMATICA'S BINARY EXECUTABLE CODE THAT RELX COPIED IN ITS NORMAL OPERATION IS DIFFERENT FROM THE COMPUTER SOURCE CODE FOR WHICH INFORMATICA REGISTERED COPYRIGHTS

70. In order to understand the distinction between computer source code and binary executable code, it's first helpful to understand that computer processors are designed to execute long sequences of relatively simple processor instructions in order to make a computer function. Although it would be possible for a human programmer to write down in a file such sequences of processor instructions for later execution on a particular processor, it would be extremely tedious

and labor-intensive to do so, and processor instructions written for certain types of processors would not work for other types of processors. Therefore, programmers instead write computer source code in human-readable computer programming languages such as C, C++, and Java. Such computer programming languages are designed to allow a computer programmer to more easily express the intended operation to any computer processor without having to understand the specific instructions that a given computer processor is capable of executing. In order to execute a program written in a computer programming language, programmers compile the computer source code into processor instructions for the particular type of processor that will execute that program. Compilers transform human-readable computer source code created by computer programmers into binary executable processor instructions that are more difficult for a human to read.

71. As an initial matter, I understand from counsel and from Mr. Groff[48] that RELX does not possess and did not ever possess any copies of Informatica's computer source code for the Informatica Software, which means that RELX would not have been able to copy it. This is consistent with Section 2.3.1 in the MSLA, which explicitly states that Informatica's computer source code shall be escrowed instead of being provided to RELX as part of the agreement:

> 2.3.1 Informatica shall within the scope of the Support Services and on a shared cost basis with Customer and in accordance with the terms of the Escrow Agreement, escrow a copy of the Source Code for the Licensed Software Into an account with the Parties' escrow agent for the purposes of this Section 2.3.

[48] See, e.g. Groff Email: "Did you ever have access to any source code created by Informatica? <DG> As stated above, No."

A. Informatica's Certified Copyright Deposits Contain the Contents of Computer Source Code Files and do not Contain the Contents of Any Binary Executable Files

72. I have reviewed the public catalog entries available on the US Copyright Office web site[49] for the following Informatica copyright registrations (hereinafter "Copyright Registrations"): Informatica B2B Data Transformation 9.6.1, registration TX0008331442; Informatica B2B Data Exchange 9.6.1, registration TX0008331795; Informatica HParser 9.6.1, registration TX0008331933; and Informatica PowerCenter Enterprise Grid Option 9.6.1, registration TX0008331715. All Copyright Registrations list "Computer File" as the "Type of Work" in the corresponding public catalog entry except for TX0008331715, which lists "Text" as the "Type of Work" instead. None of the Copyright Registrations contain either computer source code or executable binary code in their public catalog entries. In addition, I have reviewed the certified copyright deposits (hereinafter "Copyright Deposits") that were submitted to the US Copyright Office for the Copyright Registrations as provided to me by counsel.[50] All of the Copyright Deposits contain the contents of only several computer source code files and do not contain the contents of any binary executable files. The details of the contents of the Copyright Deposits follow.

73. The Copyright Deposit for Informatica B2B Data Transformation 9.6.1, registration TX0008331442, comprises 23 pages of computer source code. The computer source code is written in C and C++ and contains method definitions for the AnchorInserter class, the CAnimWnd class, the AssistantContentDlg class, and the DummyValidationReporter class. It also contains definitions for various functions including delete_directory and find_file, among others. The Copyright Deposit for Informatica B2B Data Transformation 9.6.1, registration

[49] See, e.g., http://cocatalog.loc.gov/cgi-bin/Pwebrecon.cgi?DB=local&PAGE=First.

[50] These have been produced as Bates-numbered documents Expert Discovery000014-Expert Discovery000017.

TX0008331442, includes the contents of several computer source code files and does not include the contents of any binary executable files.

74. The Copyright Deposit for Informatica B2B Data Exchange 9.6.1, registration TX000833179, comprises 25 pages of computer source code. The computer source code is written in Java and contains definitions for the DXAccountAttributeProxy class, the DXAccountProxy class, the DXBaseAttributeProxy class, the DXContactInformationProxy class, the DXEventAttributeProxy class, the DXExceptionProxy class, the MoxPCSFSession class, the OblixAuthenticationProcessingFilter class, the OblixCredentialsResolver class, the OblixLoginModule class, the Pwc86_95ClientFactory class, and the Pwc96ClientFactory class. The Copyright Deposit for Informatica B2B Data Exchange 9.6.1, registration TX000833179, includes the contents of several computer source code files and does not include the contents of any binary executable files.

75. The Copyright Deposit for Informatica HParser 9.6.1, registration TX0008331933, comprises 23 pages of computer source code. The computer source code is written in Java and contains definitions for the AddressBookFactory class, the AddressBookProtos class, the Person class, the PhoneNumber class, the Builder class, the XMLOutputFormat class, the XMLTestSupport class, the XmlWritable class, and the XMLWriteableTest class. The Copyright Deposit for HParser 9.6.1, registration TX0008331933, includes the contents of several computer source code files and does not include the contents of any binary executable files.

76. The Copyright Deposit for Informatica PowerCenter Enterprise Grid Option 9.6.1, registration TX0008331715, comprises 21 pages of computer source code. The computer source code is written in C++ and contains method definitions for the SRemotedLBProcessRun class, the

SLBAProcessRun class, the SGridEntryBase class, the SGridConnection class, the SGridInboundConnection class, the SGridServerName class, the SGridConnectionContext class, the SGridConnectionDeleteJob class, the SGridConnectionAttachCallbackJob class, the SGridConnectionDetachCallbackJob class, and the SNLoadBalancer class. The Copyright Deposit for Informatica PowerCenter Enterprise Grid Option 9.6.1, registration TX0008331715, includes the contents of several computer source code files and does not include the contents of any binary executable files.

B. RELX did not Copy Informatica's Copyrighted Computer Source Code During the Execution of the Informatica Software on RELX Systems

77. Informatica alleges in paragraph 24 of the Billingsley Declaration and paragraph six of the Mishra Declaration that RELX made unlicensed copies of the Informatica Software by copying the Informatica Software into computer memory, which is also known as random access memory ("RAM"). The Informatica Software that RELX copied in this way comprises only binary executable code intended for direct execution on a computer processor and not human-readable computer source code, for which Informatica has registered copyrights. Binary executable code is stored on a computer's hard drive or other persistent storage and when executed by the computer's processor, portions of the binary executable code are copied as necessary into RAM so that it can be more quickly accessed by the processor. At no point in this process of computer program execution is the binary executable code transformed back into human-readable computer source code.

C. RELX did not Copy Informatica's Copyrighted Computer Source Code During the Upgrade of the Informatica Software in 2015

78. Informatica alleges in paragraph six of the Mishra Declaration that RELX made unlicensed copies of the Informatica Software by copying Informatica Software into read only

memory ("ROM") as well as persistent storage. The Informatica Software that RELX copied into persistent storage comprises only binary executable code intended for direct execution on a computer processor and not human-readable computer source code, for which Informatica has registered copyrights. The Mishra Declaration states in paragraph six that when the Informatica Software was upgraded to version 9.6.1 in 2015, Informatica Software was copied into the "read only memory" of RELX's computer systems. As an initial matter, I disagree with the statements in the Mishra Declaration that any Informatica Software whatsoever was copied into the "read only memory" of the RELX computer systems for the simple reason that "read only memory" cannot be written to and can only be read from; that is what "read only" means.

79. I have also reviewed the Upgrade Scripts that were used to upgrade the Informatica Software to version 9.6.1, and those scripts show no indication that during the upgrade process any of Informatica's computer source code, as distinct from Informatica's binary executable code, was copied into persistent computer storage. For example, none of the Upgrade Scripts contain any indication that Informatica's computer source code was compiled by the Upgrade Scripts into binary executable code. Binary executable code is the form the software must be in to actually execute on a computer processor, and the Upgrade Scripts indicate that they upgraded the RELX computer systems using Informatica's binary executable code that had already been compiled. As one example, PC Upgrade_First_Node_Prod.sh and Upgrade_Each_Node_Prod.sh both make reference in their initial comments to using "rsync" to update the "/opt/inform/Informatica/9.6.1" directory before running the Upgrade Scripts.[51] "rsync" is a Unix command line command that is used to verifiably copy files from one location on a computer to another location on the same or a different computer, and the /opt directory in

[51] Mr. Groff also confirmed in the Groff Email that "rsync" was executed to perform this copy.

Unix and Linux computer systems is the typical directory in which to store compiled software installation packages, not source code.[52] This indicates that before the Upgrade Scripts upgraded the Informatica Software during the execution of the Upgrade Scripts, the binary executable code of the Informatica Software was already copied to the appropriate place on the RELX computer systems. As another example, Install.sh verifies that the currently installed version of the "java" binary executable, which itself is used to execute binary executable Java code, is supported. Install.sh does not also verify that the "javac" binary executable, which is used to compile Java source code into executable binary code, is supported.

X. FURTHER WORK

80. I may prepare visual aids to demonstrate various aspects of my testimony at trial. I may also review additional materials produced or otherwise prepared by parties or experts in this case. I reserve the right to supplement my opinions if Informatica or Informatica's expert(s) or other parties related to this case produce new materials or source code and/or if court decisions require me to clarify or provide additional support for my opinions. I also reserve the right to supplement this Report based on additional documents or information, including but not limited to additional deposition testimony, additional claims asserted by Informatica or RELX, or if Informatica's expert(s) offer any opinions.

Date: May 23, 2018

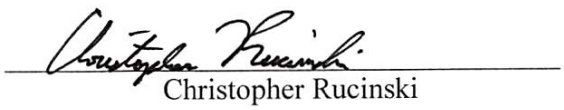

Christopher Rucinski

[52] See, e.g., https://refspecs.linuxfoundation.org/FHS_3.0/fhs/ch03s13.html, attached to this Report as Exhibit K.